Blind Statistical Steganalysis of Additive Steganography Using Wavelet Higher Order Statistics

Taras Holotyak[1], Jessica Fridrich[1], Sviatoslav Voloshynovskiy[2]

[1] Department of Electrical and Computer Engineering, State University of New York at Binghamton, Binghamton, NY, 13902-6000, USA
{holotyak, fridrich}@binghamton.edu
[2] Department of Computer Science, University of Geneva, 24 rue Général Dufour, 1211 Geneva 4, Switzerland
svolos@cui.unige.ch

Abstract. In the paper, we advocate a new approach to blind steganalysis based on classifying higher-order statistical features derived from an estimation of the stego signal in the wavelet domain. The proposed approach is flexible and enables reliable detection of presence of stego messages embedded using a wide range of steganographic methods that include ±1 embedding (LSB matching), LSB embedding, Stochastic Modulation, and others. The method is tested on raw, never compressed digital camera images, scans of photographs and films, as well as preprocessed (JPEG compressed, downsampled) images. The performance of this method is compared to the current state-of-the-art steganalytic methods for additive steganography.

1 Introduction

The goal of steganography is stealth communication – to hide a secret message in innocuous looking cover objects, such as digital media files, so that the presence of the hidden data is statistically undetectable. Furthermore, in the passive warden scenario, robustness of the hidden data to distortion is not required. Steganalysis is the art of detecting the presence of hidden messages – attacking steganography. Steganalytic methods can be roughly divided into targeted methods [1–3], in which one assumes the full knowledge of the embedding algorithm, and blind methods [4–10] in which no knowledge of the hiding algorithm is assumed. Instead, blind methods try to characterize typical cover objects in some feature space and detect stego objects by measuring how much a given image is compatible with the cluster of cover images. Classifiers trained on a large number of natural images are usually used for this purpose. This strategy is also taken in this paper.

Some steganographic embedding methods, e.g., LSB embedding, have been very successfully attacked in the past [1–3], while other embedding paradigms, such as embedding by adding noise (Stochastic Modulation [11], ±1 embedding [12] also known as LSB matching [13,14]) are much more difficult to detect.

In this paper, we address steganographic methods in which the stego image y is obtained from the cover image x by adding a low-amplitude stego noise signal $g(s, K)$ that may depend on secret stego key K and the secret data s

$$y = f(s, K) = x + g(s, K). \quad (1)$$

The function $g(.)$ represents the embedding rule. It is shown in the section dealing with experiments that the proposed detection method works quite well even for steganographic schemes in which the stego signal g exhibits some dependence on the cover image x, as is the case, for example, for LSB embedding.

The new detection technique is a blind steganalytic method for digital images in which the features are calculated from an estimation of the stego signal obtained from the stego object y in the wavelet domain. Obviously, features calculated from the estimated stego signal are more sensitive to embedding. A linear classifier is then trained on a database of images to construct a blind steganalyzer. In our experiments, we evaluate the performance of the steganalyzer on several different images classes focusing on images that are known to be difficult for steganalysis, such as raw scans, never compressed digital camera images, and grayscale images. In Section 2, we start by briefly discussing previous approaches to blind steganalysis and then explain the details of the proposed method. In Section 3, we give experimental results and compare the performance with some previously proposed methods. The paper is concluded in Section 4.

2 Blind Steganalysis

2.1 Prior Art

The first blind classifier was proposed by Memon et al. [4] in which the authors used image quality measures as features to build a classifier distinguishing between cover and stego images. Farid [6] developed a popular blind steganalyzer based on higher-order statistics (mean, variance, skewness, and kurtosis) of wavelet domain representation of stego image and the error between the logarithm of actual coefficients and the logarithm of coefficient magnitudes estimated from a globally optimal linear predictor. Kharazi et al. [5] built a blind stego classifier based on binary similarity measures as features. In another approach, Harmsen et al. [8] described a method that exploits properties of the center of mass of the Fourier transform of the image histogram. The blind steganalyzer presented in [9] is based on rate-distortion curves. The work of Fridrich [10] focused on blind steganalysis of JPEG images.

While most blind steganalyzers are relatively successful in detecting steganographic methods for JPEG images, such as OutGuess, J-Steg, and others, their performance for spatial domain steganography is less satisfactory. Also, for spatial domain steganography the performance is highly dependent on the type of imagery presented for training and testing. It is an established fact nowadays [3,13,14] that detection of steganography is significantly more difficult for scans, never compressed images, and grayscale images, and notably easier for images that were previously processed using JPEG [3,15] or for color images [8,19]. The best results for detection

of embedding by noise adding (for LSB matching or ±1 embedding) in the spatial domain were reported by Ker [13,14]. This is why in this paper we compare our results with this method and we do so for the same database of images.

2.2 The Proposed Method

The general structure of the proposed steganalysis method consists of 3 main stages: 1) stego signal estimation; 2) feature extraction; 3) classification. We now describe all three stages, focusing more attention on the feature extraction problem.

2.2.1 Stego Signal Estimation

The first step of the proposed detection is estimation of the stego signal $g(s, K)$. For this purpose, we use wavelet transform based on orthonormal Db8 basis. Although the choice of decomposition type may have a significant impact on the quality of the estimation, investigation of this issue is beyond of the scope of this paper. The purpose of stego message estimation is to remove or at least minimize the impact of the cover image and to obtain a signal that is more sensitive to embedding changes. The estimation is based on modeling the stego image in the wavelet domain as a mixture of cover image (represented with a non-stationary Gaussian model) and stego signal (modeled with a stationary Gaussian model $N(0, \sigma_S^2)$ with $\sigma_S = 0.5$). This stego signal model was selected because it enables closed form solution. We acknowledge, though, that the stego signal is in general non-Gaussian, especially in the finest decomposition levels. This estimation method has been applied in the past for steganalysis in [16] and was originally described for image denoising in [17,18].

2.2.2 Feature Extraction

Selection of the appropriate feature space plays an important role in building the stego classifier. The features must be sensitive to embedding modifications and be rather insensitive to the image content. The estimate of the stego signal from the previous section satisfies both requirements. Following the minimum description length principle, we also prefer smaller sets of features to larger feature sets.

While the energy of natural images is concentrated mainly in coarse levels, the energy of the stego signal for most steganographic techniques based on noise adding is distributed practically uniformly across the wavelet levels. We have verified this statement for the ±1 embedding paradigm (Fig. 1a). This leads to different SNRs (stego signal to cover image ratios) in different decomposition levels. Because the probability of correct stego signal detection (correct detection of stego signal presence or absence) is proportional to the SNR, $P_{corr} \propto$ SNR, from Fig. 1a we can estimate P_{corr} and $P_{err}=1-P_{corr}$ across different levels (Fig. 1b).

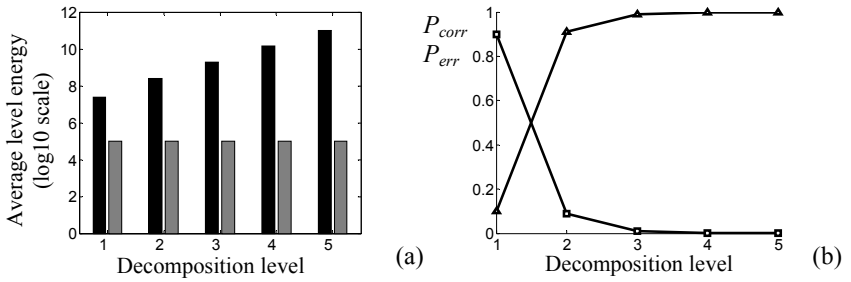

Fig. 1. a) Energy distribution for stego signal (gray) and cover image (black) in decomposition levels (averaged over 3 orientations); b) probability of correct stego signal detection (squares) and probability of error of stego signal detection (triangles) in decomposition levels.

Fig. 1 shows that the stego signal can be most accurately estimated in the finest decomposition level 1, while in higher levels the energy difference, together with a smaller number of samples, cause an increase in P_{err}. Thus, in this paper we only utilize the first decomposition level (contrary to [6] where higher decomposition levels are used). This reduces the dimensionality of the feature space and simplifies the construction of the classifier.

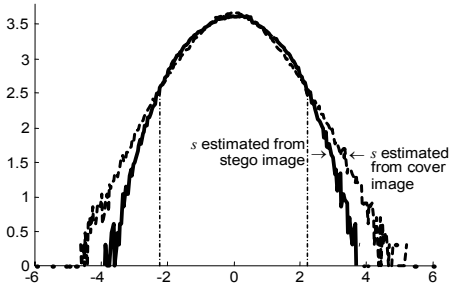

Fig. 2. Histograms of the stego signal estimation obtained from stego (solid) and cover (dashed) images (for better visualization, the y axis is in $\log_{10}$ scale).

Fig. 2 shows the histograms of the first level decomposition of the estimated stego signal extracted from the stego image $y = x + s$ and from the cover image x for the case of ±1 embedding. We see that the histograms significantly differ in their tails (high amplitude or most significant subband coefficients) located outside of the marked central zone. This means that the estimated stego signal is indeed sensitive to embedding, which enables us to separate cover and stego images.

To distinguish between the PDFs of the estimated stego signal for cover and stego images denoted p_0 and p_1, respectively, we could employ one of the existing distance measures, e.g., the Kullback-Leibler distance (KLD). We did not choose this approach due to potential problems with selecting the bin width for numerical integration necessary to calculate the KLD. Another possibility is to parametrize the PDFs

using a model and simplify the KLD so that it operates only with the models' parameters. However, finding an appropriate model that can accurately represent the stego signal for different embedding methods and stego signal parameters appears to be a difficult task. To overcome this problem, we map p_0 and p_1 to another domain (a parameter space) and use the Euclidean distance in the parameter space as the distance measure. This approach avoids the need for an explicit form for the PDFs.

Applying real-valued exponential kernel and Taylor series expansion of exp(tx), a polynomial description of the stego signal estimation in transform domain is obtained

$$M(x) = \int_{-\infty}^{\infty} p(t)\exp(tx)\,dt = \int_{-\infty}^{\infty} p(t)\left[1 + tx + \frac{t^2x^2}{2!} + \cdots + Q_n(t)\right]dt = 1 + m_1 x + \frac{m_2}{2!}x^2 + \cdots + R_n(x), \tag{2}$$

where m_i is the i^{th} moment and Q_n and R_n are residuals approaching zero with increasing n and fixed x. The function $M(x)$ is known in statistics as the moment generating function. We take the moments as features for the blind steganalyzer.

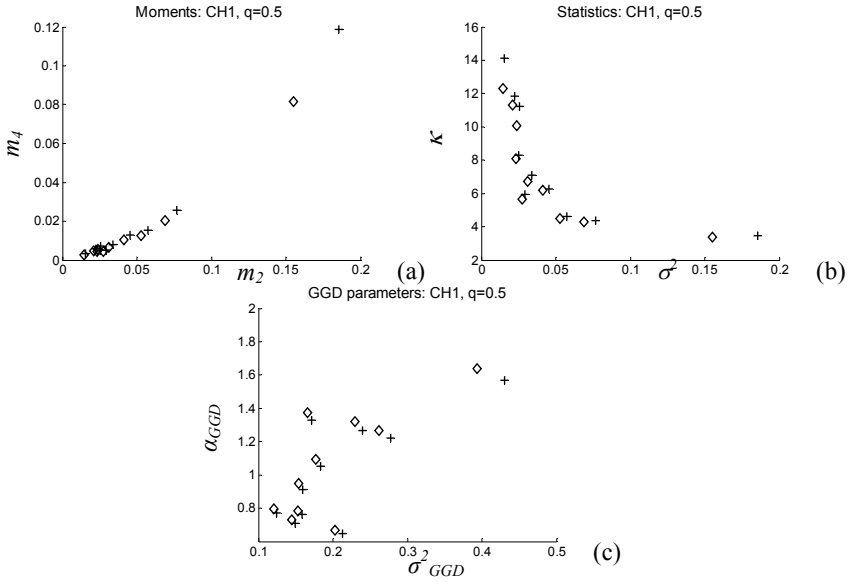

Fig. 3. Parameters of the estimated stego signal for cover (plus) and stego images (diamond) calculated for the CH1 subband for the set of 10 cover images and 10 stego images embedded with 50% relative message length using ±1 embedding: a) 2nd and 4th moments, b) variance and kurtosis, c) σ^2_{GGD} and α_{GGD} of the GGD model.

At this point, we would like to point out that the prediction errors used in [6] are also coming from some heuristic estimator of the stego signal. Denoting the stego

image and the estimated cover image in the wavelet domain as $\hat{S}$ and $\hat{X}$, the authors calculate their features from Y and $\log\left(|Y|/|\hat{X}|\right)$ as opposed to our approach in which $\hat{S} = Y - \hat{X}$ is considered. Because the PDF of $\hat{S}$ is symmetrical, we do not need to consider odd moments and thus reduce the dimensionality of the feature space by one half. Contrasting this with the approach in [6], both odd and even moments must be considered because the PDF of $\log\left(|Y|/|\hat{X}|\right)$ is typically asymmetric.

In Fig. 3, we show for a set of 10 cover and stego images (embedded with 50% relative message length using ±1 embedding) the m_2 and m_4 moments, variance σ^2 and kurtosis κ (normalized m_2 and m_4 moments), and σ^2_{GGD} and α_{GGD} – parameters of the generalized Gaussian distribution (GGD) model, which is widely used to model the distribution of wavelet subband coefficients. The GGD parameters are a unique function of the moments $\{\sigma_{GGD}, \alpha_{GGD}\} = f(m_2, m_4)$.

Clearly, the first two even moments or their transformed forms do not allow separation between cover and stego images. We need quantities that would quantify the differences in the *tails* of the distribution (see Fig 2b). Fig. 4 shows the advantages of using higher order moments to capture this difference. While it is not possible to distinguish between stego message estimations obtained from cover and stego images based solely on m_2 or m_4, the classes become practically separable when higher order moments m_8, m_{12}, or m_{20} are considered.

To investigate which moments are best suited for classification, in Fig. 5 for each moment we plot the probability of correct classification using simple Fisher linear discriminator. We see that the probability of correct classification of cover images decreases with the increasing order of moments while the accuracy of detection of stego images increases up to roughly 18 and then starts decreasing as well. We attribute the presence of this local maximum to the following two facts:

a) the low order moments are more influenced by the cover image than higher order moments because the stego signal estimation is not perfect;
b) very high order moments are too sensitive to different realizations of the stego signal and insufficient statistics.

Fig. 5b shows the product of probabilities of correct classification of both cover and stego images. We use it to determine the number of moments that should be used according the min-max criterion in the sense that the maximal probability of correct classification (highest classification accuracy) is obtained using the minimal number of features. This criterion was satisfied by moment #12. At the same time, other moments also contain information useful for classification. Therefore, we propose to use all moments for which the joint probability of successful classification is better that random guessing, i.e., more than 0.5. Based on this reasoning, we chose for further analysis the 33 moments of $\hat{S}$ normalized by its standard deviation $\sqrt{m_2}$, i.e., $m'_i = m_i / \sqrt{m_2^i}$, where $i = 4, 6, \ldots$, and 24, calculated for three different orientations (Fig. 5b). This amounts to a 33-dimensional feature vector.

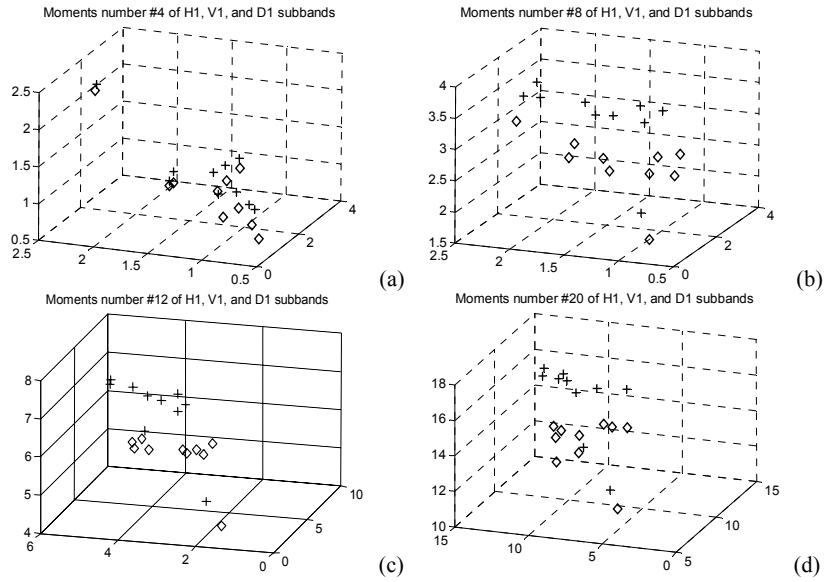

Fig. 4. Feature space formed using m'_4 (a), m'_8 (b), m'_{12} (c), and m'_{20} (d).

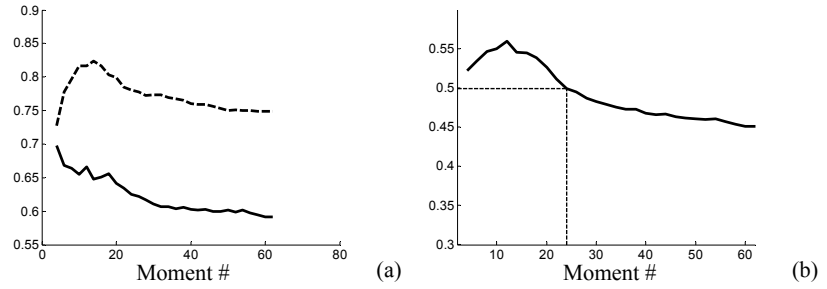

Fig. 5. Impact of the moment order on classification accuracy: a) probability of correct classification of cover images (solid), probability of correct classification of stego images (dashed); b) joint probability of successful steganalysis vs. order of moments averaged over 500 raw images embedded using ±1 embedding and LSB embedding with relative message length in range 0.125–0.5.

2.2.3 Classification

Naturally, the features defined in the previous section exhibit mutual dependencies that can be removed using Karhunen-Loeve transform that optimally decorrelates the features in the mean-square error sense. This approach is also called principal component analysis. The transform also enables reducing the dimensionality of the feature

space. Our analysis shows that the 33-dimensional feature space can be accurately represented in a 4–5 dimensional decorrelated feature space.

To investigate the practical performance of the proposed steganalyzer and the impact of feature selection, we used one of the simplest classifiers – the Fisher linear discriminator (FLD). In FLD, the feature space is projected on a one-dimensional space, where various decision rules can be applied for determining the classification thresholds. We acknowledge that potentially better classification results could be obtained using more sophisticated, preferable, multidimensional classifiers.

3. Numerical Results

It is a well established fact that never compressed raw images pose the biggest challenge for steganalysis [3,13,14]. Especially scans of film and photographs are difficult for steganalysis due to a high level of noise that interferes with the stego signal. It is also more difficult to detect stego signals in grayscale images as it is not possible to use the relationship between color channels as in [8,19]. In fact, the performance of [8] decreases quite noticeably when color information is ignored. Also, for some methods color information is essential for the method to work [19]. A fundamental reason why color helps to such an extent is that the 3D color histogram is much sparsely populated (due to the high number of colors) than a histogram of a grayscale image. Consequently, embedding by noise adding leaves more noticeable artifacts in the color histogram than in the grayscale histogram. Finally, in general, steganalysis is easiest for images that were previously compressed using JPEG because of their lower energy of high-frequency noise removed during compression. We also note that for decompressed JPEGs, methods similar to JPEG compatibility steganalysis are possible and can be very accurate. We decided to test the proposed approach on the following two classes of raw images.

Set #1 consists of 2567 raw, never compressed color images of different dimensions (all larger than 1 megapixel), some stored in the 48-bit TIFF format, some in 24-bit BMP format acquired by 22 different digital cameras ranging from low-cost cameras to semiprofessional cameras. Part of this database was downloaded from [20]. For our experiments, we converted the color images to grayscale and, where it was required, decreased the color depth to 8 bits.

Set #2 includes high resolution (1500×2100 pixels) 32 bit CMYK color TIFF images (2375 images) from [21]. To be able to compare our results to [14], similar as the authors we converted all color images to grayscale and applied bicubic downsampling.

The receiver operating characteristics (ROCs) for Set #1 and #2 together with an example of an ROC for a single image source (one digital camera Canon PowerShot S40) are presented in Fig. 6. The embedding method tested was ±1 embedding also called LSB matching. We also report some quantitative characteristic commonly used in the literature, such as the false positive rate at 50% stego detection rate (Table 1).

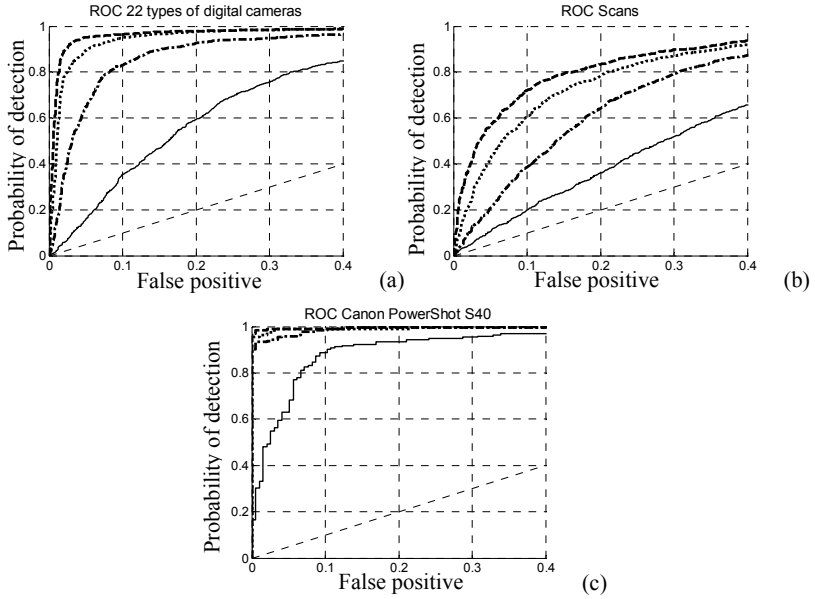

Fig. 6. ROCs for ±1 embedding for Set #1 (a), #2 (b) and a single camera (Canon PowerShot S40) (c) with different embedding capacity: solid = 0.25 bits per pixel (bpp), dash-dotted = 0.5 bpp, dotted = 0.75 bbp, dashed = 1 bpp.

Table 1. False positives at 50% correct detection of stego images.

Image source	Embedding ratio (bpp)			
	0.25	0.5	0.75	1.0
Canon Power Shot S40	2.56%	0%	0%	0%
Set #1	16.24%	3.12%	1.17%	0.58%
Set #2	28.63%	13.98%	6.57%	3.45%

We can see that the detection of embedding in scans (Set #2) is less reliable than for digital camera images. Again, this is due to the high level of noise in scans that interferes with the stego signal. Fundamentally, it appears to be difficult to distinguish between the stego noise and noise naturally present in scans. As the same Set #2 was used in experiments in [14], we can make direct comparison of performance. The false positives for the proposed scheme for detection rate 50% and 80% were 3.45% and 16.25% compared to about 7% and 27% reported in [14].

Contrasting Fig. 6a with Fig. 6c and Table 1, we see that for a homogenous image source (one camera) it is possible to detect ±1 embedding rather reliably even at low embedding rates (for grayscale never compressed images). Similar results were obtained for the other 21 cameras.

Just for comparison, we used the proposed steganalysis method to detect the LSB embedding in raw images (test was done for 100 512×512 images cropped from Kodak DC290 set) and ±1 embedding in decompressed JPEG cover images (ROC was built for 100 512×512 images cropped from Canon Power Shot G2 set that were previously JPEG compressed with quality factor 90%). The results (Fig. 7) confirm our statement made above that detection of steganography in previously compressed images is significantly easier – the detection is nearly perfect even for 0.25 bits per bit, which quite an impressive result for a blind scheme.

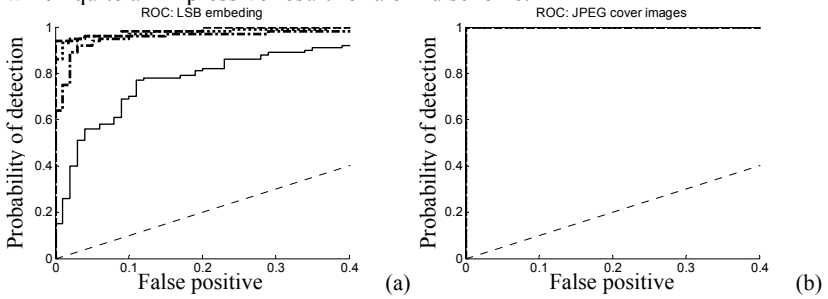

Fig. 7. ROCs for LSB embedding in raw images (a) and ±1 embedding in previously compressed JPEG images (b) with different capacity: solid = 0.25 bpp, dash-dotted = 0.5 bpp, dotted = 0.75 bbp, dashed = 1 bpp.

The last set of results is related to the classification ability of the proposed method. For this purpose, the Monte-Carlo approach was applied to Set #1 and #2. The probability of correct detection of non-embedded images P_0 and correct detection of embedded images P_1 were obtained by averaging over 100 different random choices of training subsets with 2:8 ratio between the size of the training and testing subsets for each Set (Table 2). The classification was done using the FLD with Bayes and Neyman-Pearson ($P_0 = 0.8$) decision rules.

Table 2. Correct classification of cover P_0 and stego images P_1 using FLD

Embedding ratio (bpp)	P_0/P_1 Bayes decision rule		P_0/P_1 Neyman-Pearson decision rules	
	Set #1	Set #2	Set #1	Set #2
0.25	0.71 / 0.72	0.61 / 0.62	0.79 / 0.58	0.79 / 0.37
0.50	0.86 / 0.87	0.72 / 0.73	0.79 / 0.92	0.79 / 0.63
0.75	0.92 / 0.92	0.78 / 0.79	0.79 / 0.97	0.79 / 0.78
1.00	0.94 / 0.94	0.80 / 0.81	0.79 / 0.98	0.79 / 0.82

Conclusions

In this paper, we describe a new approach to blind steganalysis realized using a linear classifier with features calculated from higher-order statistical moments of PDF of the estimated stego signal in the finest wavelet level. Because the features are calculated

from the estimated stego signal, they are more sensitive to steganographic modifications while suppressing the influence of the cover image.

The proposed method is tested on various classes of images that are known to pose problems for steganalysis – never compressed raw images from digital camera and grayscale uncompressed film scans. We test the methodology on the ±1 embedding paradigm and LSB embedding. On raw grayscale digital camera images for ±1 embedding, we obtained reliable detection results for message lengths above 0.5 bits per pixel. For images coming from a homogenous source, such as raw grayscale images obtained using a single camera, relatively reliable detection is even possible at the embedding rate of 0.25 bits per pixel (for ±1 embedding).

The detection performance on decompressed JPEGs embedded with ±1 embedding showed why such images are the easiest class for steganalysis – the detection was nearly perfect even for embedding rates of 0.15 bits per pixel, which is a very good result for a blind steganalysis approach.

Acknowledgment

The work on this paper was supported by Air Force Research Laboratory, Air Force Material Command, USAF, under a research grant number F30602-02-2-0093 and FA8750-04-1-0112 (T. Holotyak and J. Fridrich) and Swiss National Foundation Professorship grant No PP002-68653/1 (S. Voloshynovskiy). The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation there on. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of Air Force Research Laboratory, or the U. S. Government. Authors would like to thank to Miroslav Goljan for many useful discussions.

References

1. J. Fridrich, M. Goljan, D. Hogea, and D. Soukal: "Quantitative Steganalysis: Estimating Secret Message Length," ACM Multimedia Systems Journal, Special Issue on Multimedia Security, vol. **9**(3), pp. 288–302, 2003.
2. S. Dumitrescu, X. Wu, and Z. Wang: "Detection of LSB Steganography via Sample Pair Analysis", in: Petitcolas, F.A.P. (ed.): Information Hiding. 5th International Workshop. Lecture Notes in Computer Science, vol. 2578, Springer-Verlag New York, pp. 355–372, 2000.
3. A. Ker: "Improved Detection of LSB Steganography in Grayscale Images", in J. Fridrich (ed.): Information Hiding. 6th International Workshop. Lecture Notes in Computer Science, vol. 3200, Springer-Verlag New York, pp. 97–115, 2005.
4. I. Avcibas, N. Memon, and B. Sankur: "Steganalysis Using Image Quality Metrics", in E. Delp et al. (eds.): Proc. SPIE Electronic Imaging, Security and Watermarking of Multimedia Contents II, vol. 4314, pp. 523–531, 2001.

5. M. Kharrazi, H. T. Sencar, N. Memon: "Benchmarking Steganographic and Steganalytic Techniques", to appear in E. Delp et al. (eds.): Proc. SPIE Electronic Imaging, Security, Steganography, and Watermarking of Multimedia Contents VII, 2005.
6. H. Farid and L. Siwei: "Detecting Hidden Messages Using Higher-Order Statistics and Support Vector Machines", in: F.A.P. Petitcolas (ed.): Information Hiding. 5th International Workshop. Lecture Notes in Computer Science, vol. 2578, Springer-Verlag New York, pp. 340–354, 2002.
7. R. Tzschoppe, R. Bäuml, J. Huber, and A. Kaup: "Steganographic System Based on Higher-Order Statistics", in E. Delp et al. (eds.): Proc. SPIE Electronic Imaging, Security, Steganography, and Watermarking of Multimedia Contents V, vol. 5020, pp. 156–166, 2003.
8. J. J. Harmsen and W. A. Pearlman: "Steganalysis of Additive Noise Modelable Information Hiding," in E. Delp at al. (eds.): Proc. SPIE Electronic Imaging, Security, Steganography, and Watermarking of Multimedia Contents V, pp. 131–142, 2003.
9. M. Celik, G. Sharma, and A. Tekalp: "Universal Image Steganalysis Using Rate-Distortion Curves, in E. Delp et al. (eds.): Proc. SPIE Electronic Imaging, Security, Steganography, and Watermarking of Multimedia Contents VI, vol. 5306, pp. 467–476, 2004.
10. J. Fridrich: "Feature-Based Steganalysis for JPEG Images and its Implications for Future Design of Steganographic Schemes", in J. Fridrich (ed.): Information Hiding, 6th International Workshop. Lecture Notes in Computer Science, vol. 3200, Springer-Verlag New York, pp. 67–81, 2005.
11. J. Fridrich and M. Goljan: "Digital Image Steganography Using Stochastic Modulation," in E. Delp et al. (eds.): Proc. SPIE Electronic Imaging, Security, Steganography, and Watermarking of Multimedia Contents V, vol. 5020, pp. 191–202, 2003.
12. T. Sharp: "An Implementation of Key-Based Digital Signal Steganography," in I. Moskowitz (ed.): Information Hiding. 4th International Workshop. Lecture Notes in Computer Science, vol. 2137, Springer-Verlag New York, pp. 13–26, 2001.
13. A. Ker: "Resampling and the Detection of LSB Matching in Color Bitmaps", to appear in E. Delp et al. (eds.): Proc. SPIE Electronic Imaging, Security, Steganography, and Watermarking of Multimedia Contents VII, 2005.
14. A. Ker: "Steganalysis of LSB matching in grayscale images", to appear in IEEE Signal Processing Letters, 2005.
15. J. Fridrich, M. Goljan and R. Du: "Steganalysis Based on JPEG Compatibility", in A.G. Tescher et al. (eds.): Proc. SPIE Multimedia Systems and Applications IV, vol. 4518, pp. 275–280, 2001.
16. T. Holotyak, J. Fridrich, and D. Soukal: "Stochastic Approach to Secret Message Length Estimation in $\pm k$ Embedding Steganography", to appear in E. Delp et al. (eds.): Proc. SPIE Electronic Imaging, Security, Steganography, and Watermarking of Multimedia Contents VII, 2005.
17. S. M. LoPresto, K. Ramchandran, and M. T. Orchard: "Image Coding Based on Mixture Modeling of Wavelet Coefficients and a Fast Estimation-Quantization Framework", in Proc. Data Compression Conf., pp. 221–230, March 1997.
18. M. K. Mihcak, I. Kozintsev, K. Ramchandran, and P. Moulin: "Low-Complexity Image Denoising Based on Statistical Modeling of Wavelet Coefficients", IEEE Signal Processing Letters, vol. **6**(12), pp. 300–303, 1999.
19. A. Westfeld, "Detecting Low Embedding Rates", in: Petitcolas, F.A.P. (ed.): Information Hiding. 5th International Workshop. Lecture Notes in Computer Science, vol. 2578, Springer-Verlag New York, pp. 324–339, 2002.
20. The Digital Forensic Image Library (DFIL), http://www.cs.dartmouth.edu/cgi-bin/cgiwrap/isg/imagedb.py.
21. NRCS Photo Gallery, http://photogallery.nrcs.usda.gov.